The first step for program insertion to take place at a broadcast head-end is to extract (by demultiplexing) the packets, identified by the PIDs, of the individual elementary bit streams that make up the program, including the bit stream carrying the program_map_table. Once these packets have been extracted, as illustrated in Fig. 8.1, program insertion can take place on an individual PID basis. If applicable, some packets may be passed through without modification. There is also the flexibility to add and drop elementary bit streams. The splicing process for each PID is described in the next section.

When program insertion takes place, the program_map_table needs to be modified to reflect the properties of the program transport stream that is being spliced in. As described in the section on its syntax, the definition of the program_map_table allows the signaling of a change in the contents of a program transport stream ahead of time. The changes in the program definition could involve a change in the number of elementary bit streams that make up the program, either by addition or removal of bit streams, change in the PIDs used for the elementary bit streams, etc...

The bit-rate of the program after splicing should have a known relationship to the bit-rate of the program before splicing, in most scenarios. Unless dynamic bit-rate allocation is possible for a program at the system multiplexer (based on instantaneous bandwidth requirements), an increase in bit-rate after splicing can cause buffer overflow. A decrease in bit rate may be handled by transmitting null packets (packets with no information) or by allocating the extra bandwidth to other programs on a dynamic basis. These capabilities depend on the implementation of the system level multiplexing function, a function that is not a part of the GA Transport specification.

Bit rate constraints may also be imposed on individual elementary bit streams for the program that is inserted, e.g., for compressed video, input and output bit rates need to be the same. In a perfect program insertion set up, the splice points for the different elementary streams in a program should be coordinated to correspond to the same instant in time in the overall program (which may not correspond to the same instant in time for each elementary bit stream) to permit seamless transition. Additional constraints on selecting the splicing points exist for particular applications such as video (i.e., VBV_delay value).

5.8.2. Basics of elementary bit stream insertion

The interface for elementary bit stream insertion is at the transport layer of the protocol. This means that bit stream insertion always takes place in units of transport packets. The primary features enabling local elementary bit stream insertion are the discontinuity_indicator field and the splice_countdown fields in the transport header. The discontinuity_indicator signals the decoder that the

PCR is changing to a new time base. This simply informs the decoder that the change in the bit stream is not due to an error in the channel, but rather is intended by the program provider. The implication for the decoder is that it should continue normal decoding, and it is the encoders responsibility to make sure that the bit stream has been constructed in a compliant manner (that is that decoders don't crash due to overflow or underflow.)

The splice_countdown field in the adaptation header is used to signal a head-end or intermediate digital switch that a subsequent packet is the point for switching in a new bit stream. The count-down is a positive number which decrements on each subsequent packet of that service. The value of "0" is the last packet in the original sequence, and the value "-1" is resident in the packet which should initiate the switch over. The count will continue to decrement for channel error resilience. The behavior is shown in Figure 5.8.2.

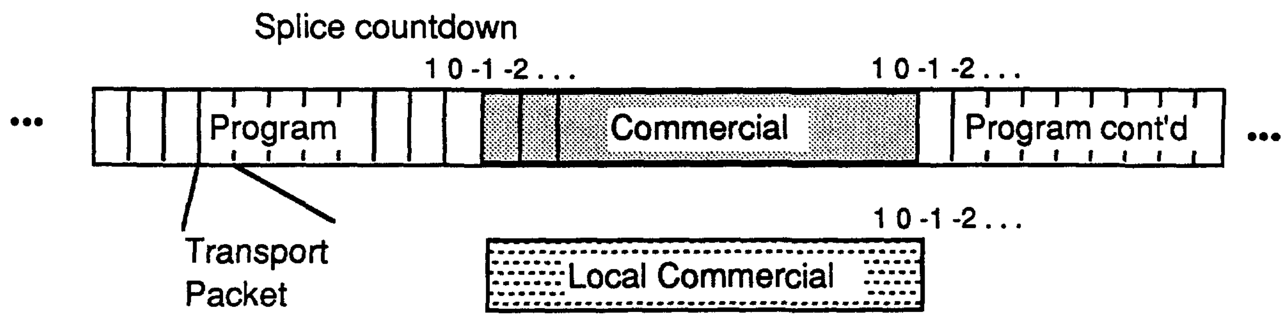

Figure 5.8.2. Local program insertion keyed on Splice countdown.

The affiliate or headend equipment sets itself up for the switch based on the descending countdown. At the "-1" point, the local commercial is inserted while the network commercial continues. At the second trailing "-1", the affiliate returns to the network feed. This technique can be used for either local programs or local commercials. It does depend on the video encoder constraining its bit generation at the splice points so that the decoder buffer does not overflow.

The GA transport encoder places some constraints on the encoding which are more stringent than the MPEG-2 requirements. The added constraint is that the PES header is followed immediately by a video access unit. This will speed acquisition. The first packet in an insertion will contain the PCR value, with the PCR discontinuity bit set to "1" to inform the decoder that a splice has occurred. The first payload in the stream will begin with a PES Header, which will have a PTS resident, so that the decoder can determine the display time immediately. Because the PES header also has the data_alignment_indicator set, the first data following the header will be the start of the video sequence layer. Consequently, the decoder has all the information available to begin decoding immediately after receiving the beginning of the spliced commercial. (In general, an MPEG-2 stream does not have these constraints imposed, and hence does not have guaranteed performance at the splice points.)

5.8.3. Restrictions

Compressed digital technology does impose some restrictions on affiliate operations which may differ from present practice. Although all present practices can be replicated by completely decoding and recoding the video, there is a desire to implement as much as possible in the compressed video domain. The following comments are made with respect to processing the compressed video.

Local commercials and network commercials will need to be strictly controlled to be the same number of packets, and the same number of frames of video. This is contrary to the present practice where local inserts may differ from the planned network inserts by several seconds.

A second restriction is that affiliate pix-in-pix and affiliate text overlays can <u>only</u> be accomplished by decoding and recoding.

5.8.4. Imperfect program insertion

It may not always be possible for the program insertion process to meet the precise requirements of a seamless splice. This could be due to several reasons including the presence of infrequent splice points in the incoming bit stream or non-availability of the hardware required for precise splicing at the network affiliate. There are two scenarios for imperfect splicing. In the first scenario the network affiliate attempts to splice in the entire program as a whole, without attempting to align each of the component elementary bit streams. In this case, the exact splicing can take place for only one of the elementary streams. Since video is the most important component of the program, perfect alignment will be obtained for video. In this case the output of the other elementary bit streams will not be presented at the output until synchronization of these bit streams is achieved. As an example, audio should be muted until its elementary bit stream is synchronized.

In the second and most uncoordinated splicing approach, the splicing takes place without any attempt at coordination with the input bit stream. In this case the video presentation process is affected around the splicing point. If the splicing takes place when the VBV level in the existing bit stream is less than it should be for perfect splice, there will be a period of time for which data is lost for the existing bit stream. In this case the decoder should freeze the last displayed frame. In the other case where VBV is fuller than expected, the decoder video data buffer may eventually overflow during the time period of the spliced in bit stream. The decoder will then have to initiate a resynchronization procedure in the middle of the program, freezing the display to the last decoded picture while this process is taking place. Note that when the process of splicing in a bit stream does not take place correctly, there will also be a disruption in service at the splice back to the original bit stream. It is the recommendation of the GA that this type of splicing be strictly prohibited, since it leads to a very noticeable interruption of service.

It is important to note that the process of facilitating frequent opportunities for splicing in a program bit stream is not within the control of the transport layer of the system. The transport only provides the mechanism of implementing the splice itself. Hence decisions on determining the possible frequency of commercial insertion should also involve the people involved in the design of the source coding algorithms for applications like video and audio.

5.9. Compatibility with other Transport Systems

The GA transport system is compatible with two of the most important alternative transport systems, namely the MPEG-2 transport stream definition, and also the ATM definition being finalized for Broadband ISDN. Furthermore, since several of the CATV (e.g., Digicipher II) and DBS systems being designed are considering use of the MPEG-2 Transport layer syntax, the degree of interoperability with such deployed systems should be quite high (possibly requiring a translation if the CATV or DBS system deploys a slightly incompatible MPEG-2 variant).

5.9.1. Interoperability with MPEG-2

In the development of the GA transport specification, the intent has never been to limit the design by the scope of the MPEG-2 systems definition. The GA system is interoperable with MPEG-2 decoders since the GA Transport is currently a constrained subset of the MPEG-2 Transport syntax. The constraints are imposed for reasons of increased performance of channel acquisition, bandwidth efficiency and decoder complexity. If, in the course of future work, the MPEG-2 standard is unable to efficiently meet the requirements of the GA system, a deviation from MPEG would be in order.

The ATV system requires definition of bit streams and services beyond the compressed video and audio services. A means of identifying such bit streams is necessary in the ATV system, but is not part of the MPEG-2 definition. There is a method of encoding such a registration descriptor when a an authority to administrate registration is identified. This identification is implemented by the registration_descriptor in the PSI stream.

5.9.2. Interoperability with ATM

The GA transport packet size is selected to ease transferring these packets in a link layer that supports Asynchronous Transfer Mode (ATM) transmission. There are several methods for mapping the Transport packet into the ATM format. Three techniques are presented, although the industry may converge to a different solution than those presented here.

5.9.2.1. ATM Cell and Transport Packet Structures

Figure 5.9.1 shows the format of an ATM cell. The cell consists of two parts: a five byte header and a forty-eight byte information field. The header, primarily significant for networking purposes, consists of the following fields:

GFC	a four bit Generic Flow Control field used to control the flow of traffic across the User Network Interface (UNI). Exact mechanisms for flow control are under investigation.
VPI	an eight bit network Virtual Path Identifier.
VCI	a sixteen bit network Virtual Circuit Identifier.
PT	a three bit Payload Type (i.e., user information type ID).
CLP	a one bit Cell Loss Priority flag (eligibility of the cell for discard by the network under congested conditions).
HEC	an eight bit Header Error Control field for ATM header error correction
AAL	ATM Adaptation Layer bytes (user specific header).

The ATM User Data Field consists of forty-eight bytes, where up to four of these bytes can be allocated to an Adaptation Layer.

Figure 5.9.2 shows the format of the Grand Alliance transport packet. A one hundred eighty-four byte packet data field (possibly including an optional and conditional adaptation field) is preceded by a four byte prefix.

5.9.2.2. Null AAL Byte ATM Cell Formation

The simplest method to form ATM cells from the Transport layer is the null AAL byte structure shown in Figure 5.9.3. The Transport packet is partitioned into forty-eight byte payloads, applied directly to the information fields of the ATM cell. The five byte ATM header is appended. Since the Transport packet length is not an integer multiple of the ATM cell payload, there will be only occasional alignment of the Transport header with the start of the ATM cell information field.

5.9.2.3. Single AAL Byte ATM Cell Formation

Alignment of the Transport packet and ATM cell is accommodated by parsing the Transport packet into forty-seven byte segments, shown in Figure 5.9.4. Four such segments will exactly encompass a Transport packet. A one byte AAL is appended, along with the five byte ATM header to fulfill the fifty-three byte ATM cell requirement. The AAL byte can carry useful information concerning the transport data within the ATM cell. It can be viewed as an adaptation field for the contained data, conveying the original position of the ATM payload within the Transport packet, for example, as well as other information. For example, ATM standards presently provide for five different AALs, such as AAL Type 1 for accommodating connection oriented constant bit rate services, and AAL Type 2 for handling connection oriented variable bit rate data services.

5.9.2.4. Dual AAL Byte ATM Cell Formation

An alternative solution to cell/packet alignment is shown in Figure 5.9.5. The transport header is discarded, and the remaining one hundred eighty-four byte payload is segmented into 46 byte increments. To these are added two AAL bytes and the five byte ATM header for each ATM packet. The idea here is that there may be a duplication in functionality in the ATM header and the link level transport header fields. A particular header field to consider for duplication of functionality is the PID. If the PID can be associated with a specific VPI and VCI used in the ATM headers of the packets carrying the data payload, and this PID mapping information can be sent to the destination terminal when the virtual path/circuit is set up (using the ATM signaling channel), it does not have to be transmitted for every GA packet. The PID can then be reconstructed at the destination (using the information transmitted at call setup), and can then be appended to the one hundred eighty-four byte payload (reconstructed from four ATM packets) to obtain the complete GA packets. Transport header information that cannot be reconstructed (e.g., adaptation field control) should be carried as a part of the two AAL bytes for each ATM cell, along with other additional information. Note that the above is only a suggested approach and does not represent a complete design.

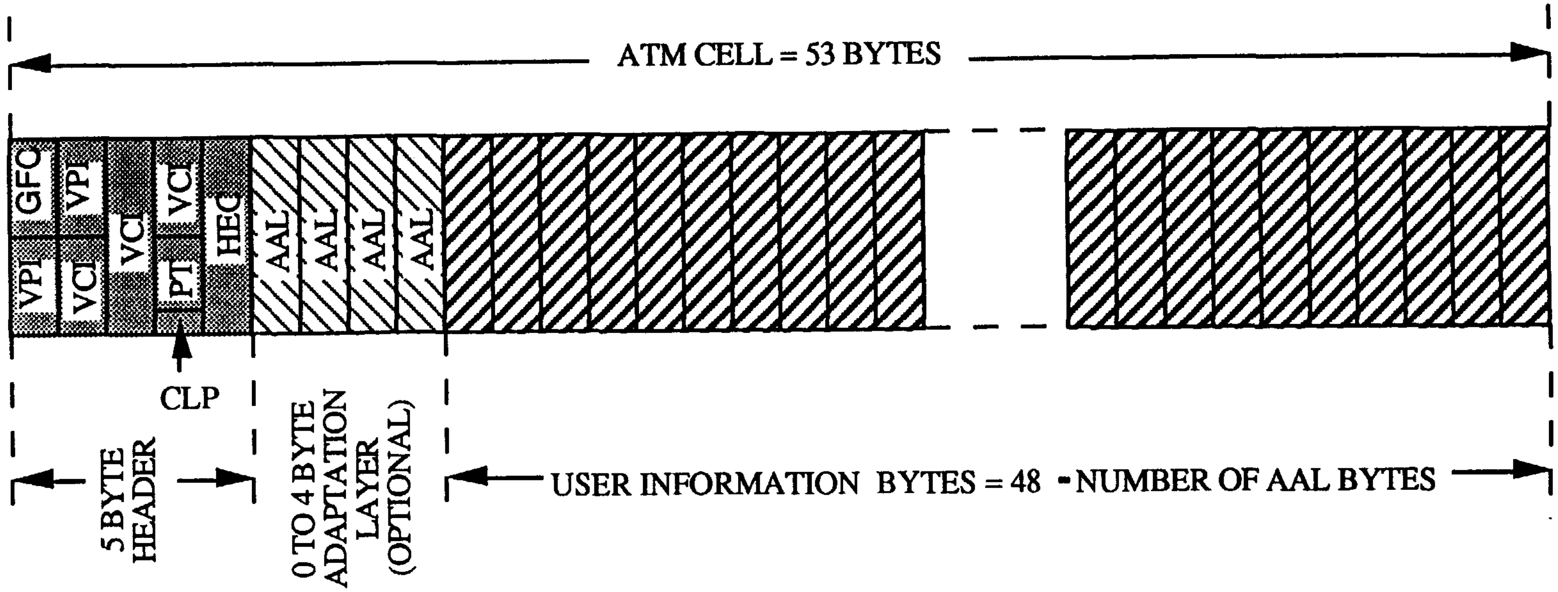

FIGURE 9.1 STRUCTURE OF THE ATM CELL

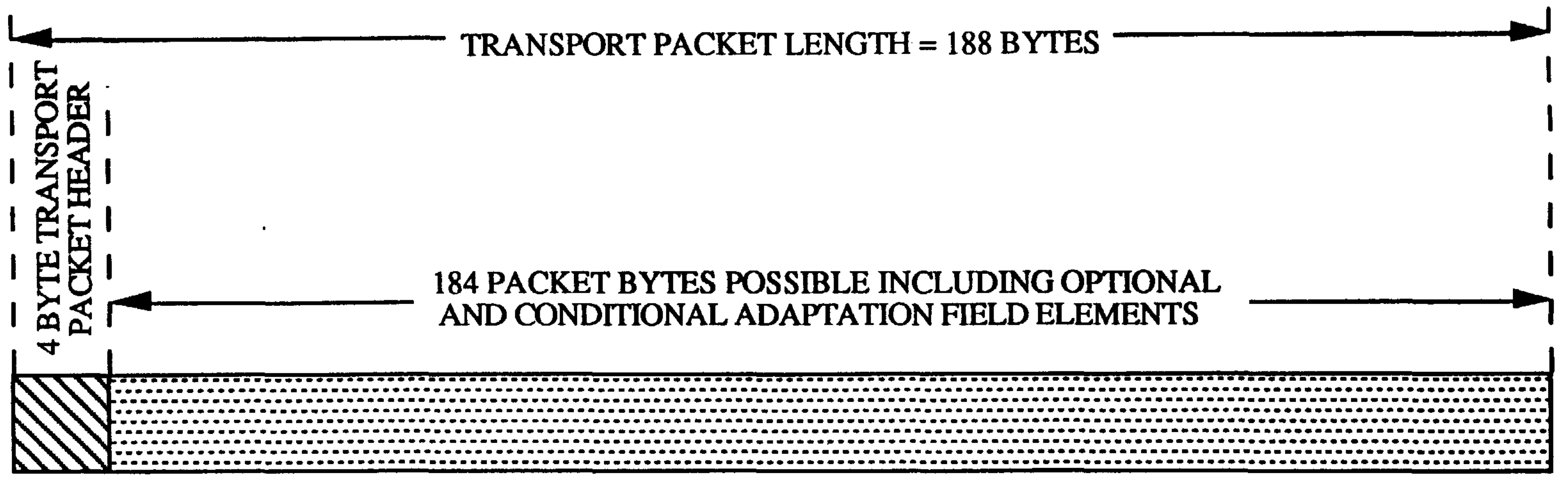

AAA 7/29/93

FIGURE 9 .2 STRUCTURE OF THE TRANSPORT PACKET

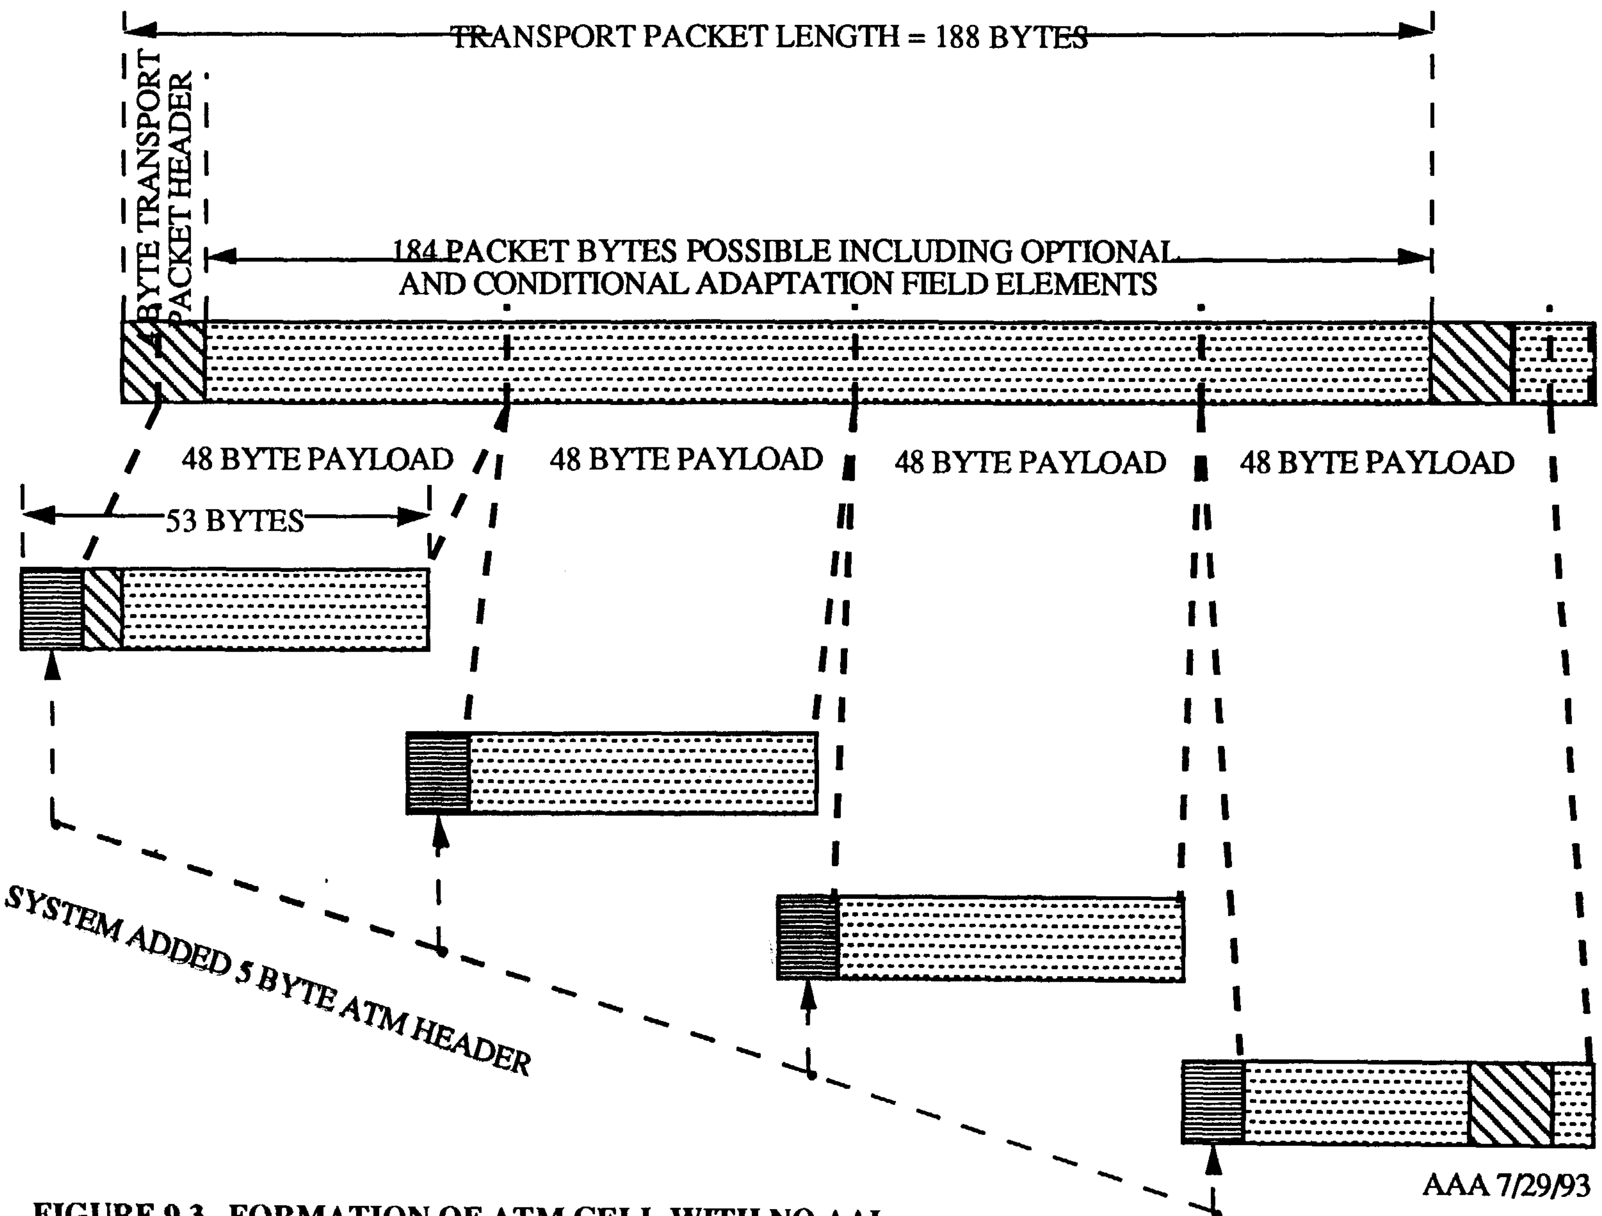

FIGURE 9.3. FORMATION OF ATM CELL WITH NO AAL BYTES

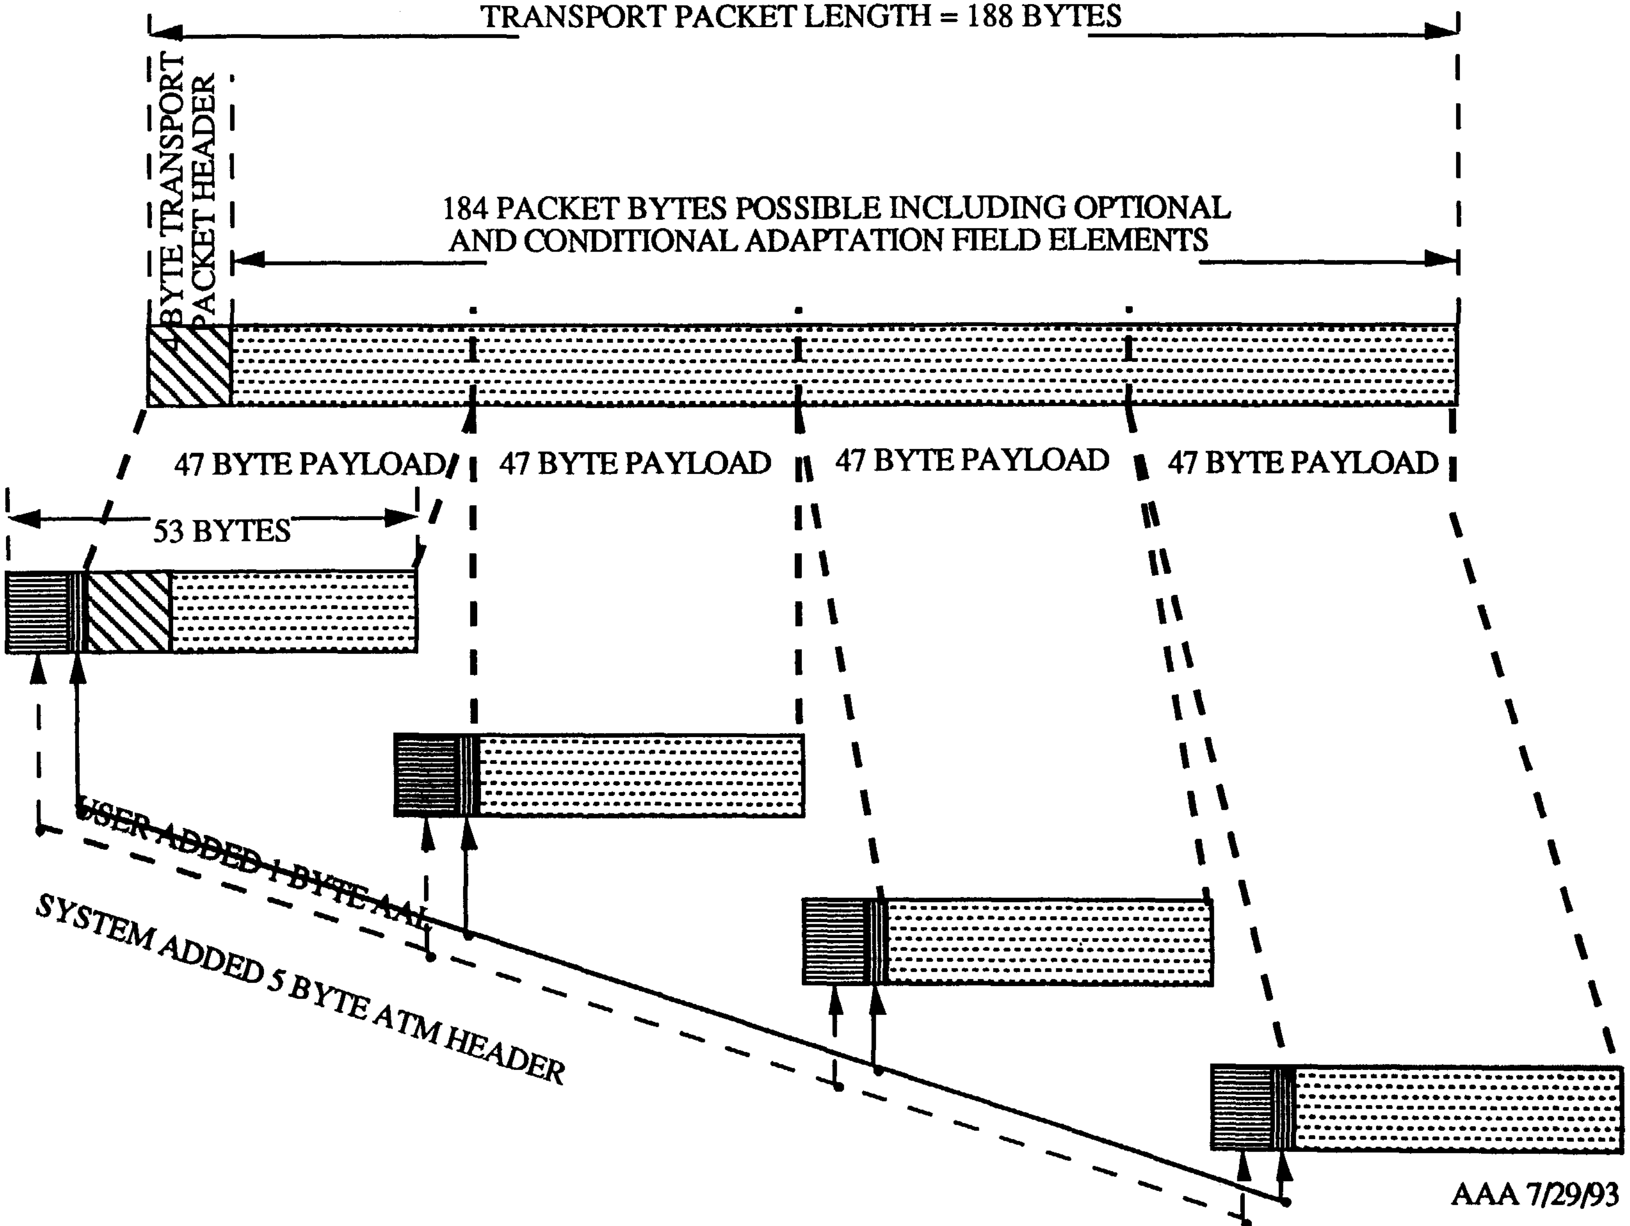

FIGURE 9.4. FORMATION OF THE ATM CELLS WITH A SINGLE AAL BYTE

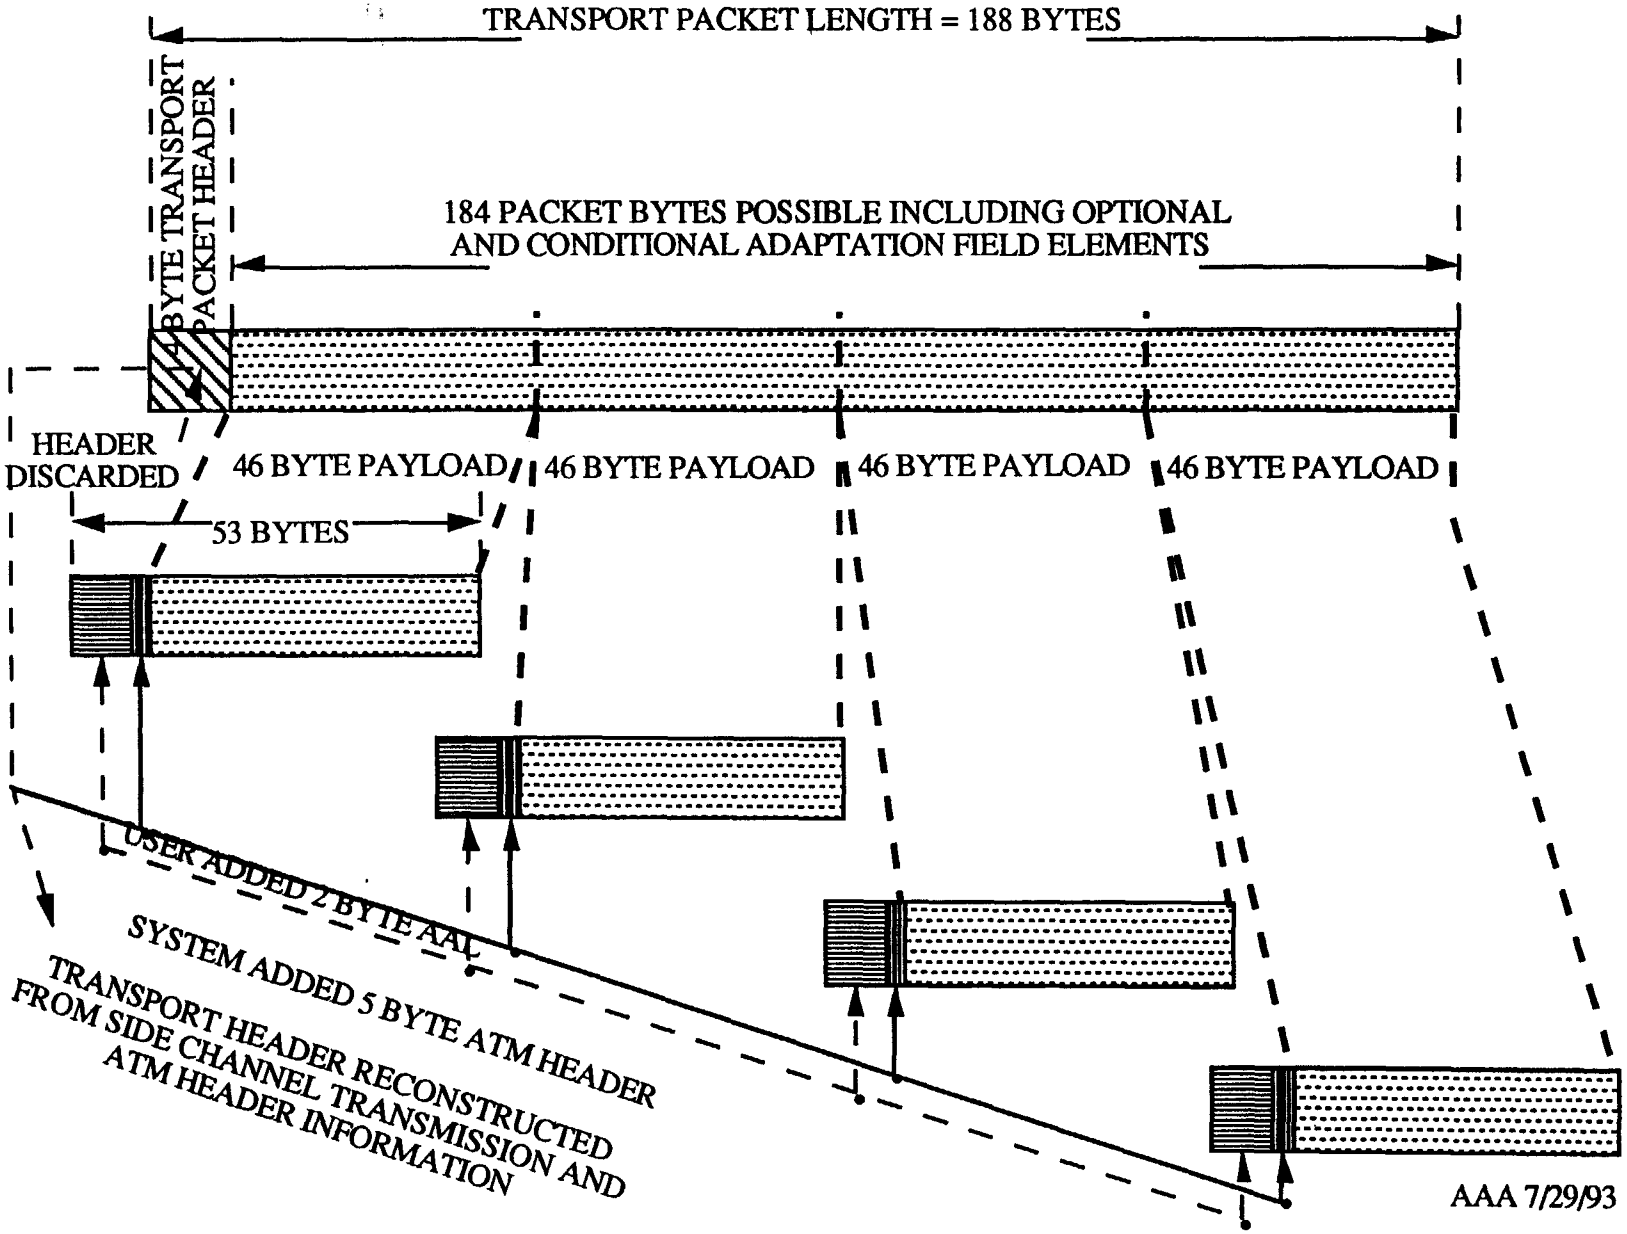

FIGURE 9.5. FORMATION OF AN ATM CELL WITH DUAL AAL BYTE

Chapter 6
TRANSMISSION SYSTEM

6.1.0 Introduction And System Overview

The Grand Alliance (G-A) vestigial sideband (VSB) digital transmission system provides the basis for a family of transmission systems suitable for data transmission over a variety of media. This family shares the same pilot, symbol rate, data frame structure, interleaving, Reed-Solomon coding, and synchronization pulses. The VSB system has two modes: a simulcast terrestrial broadcast mode, and a high data rate cable mode. The terrestrial broadcast mode, which transmits an ATV signal on currently unusable taboo NTSC channels with minimal interference to or from NTSC channels, is optimized for maximum coverage area, and supports one ATV signal in a 6 MHz channel. The high data rate cable mode, which trades off some robustness for twice the data rate, supports two ATV signals in one 6 MHz channel.

The VSB transmission systems (both modes) take advantage of a pilot, a segment sync, and a training sequence for robust acquisition and operation. The two systems also share identical carrier, sync, and clock recovery circuits, as well as phase correctors and equalizers. Additionally, both systems use the same Reed-Solomon (R-S) code for forward error correction (FEC).

In order to maximize coverage area, the terrestrial broadcast mode incorporates both an NTSC rejection filter (in the receiver) and trellis coding. In contrast to the 4-VSB system tested during the first round of ATV testing, precoding at the transmitter is incorporated in the trellis code. When the NTSC rejection filter is activated in the receiver, the trellis decoder is switched to a trellis code corresponding to the encoder trellis code concatenated with the filter.

The cable mode, on the other hand, does not have as severe an environment to work in as that of the terrestrial system. Therefore, a higher data rate is transmitted in the form of more data levels (bits/symbol). For cost considerations, no trellis coding or NTSC interference rejection filters are employed.

VSB transmission inherently requires only processing the in-phase (I) channel signal, sampled at the symbol rate, thus optimizing the receiver for low cost implementation. The decoder only requires one A/D converter and a real equalizer operating at the symbol rate of 10.76 Msamples/second.

The parameters for the two VSB transmission modes are shown in Table 6.1.

6.2.0 Terrestrial Vsb System Description

6.2.1 System Information

The VSB transmission system transmits data according to the data frame depicted in Fig 1. The frame is organized into segments each with 836 symbols. Each transmitted segment consists of a four symbol segment sync followed by 832 data plus FEC symbols. The data contained in each transmitted 208 byte segment is a 188 byte MPEG-compatible data packet with 20 R-S parity bytes.

The exact symbol rate is:

4.5/286 Mhz x 684 = 10.76 MHz

The frequency of a segment may be calculated as follows:

f_{seg}=10.76 Msymbols/sec / 836 symbols/seg = 12.87 ksymbols/sec

For terrestrial broadcast mode, each segment corresponds to one R-S correction block of 208 bytes as follows:

208 bytes/block x 8 bits x 3/2 trellis = 3 bits/symbol x 832 symbols

For the high data rate cable mode, each segment corresponds to two R-S correction blocks of 208 bytes as follows:

2 blocks x 208 bytes/block x 8 bits = 4 bits/symbol x 832 symbols

As shown in Fig 1, each data frame begins with a first data field sync segment followed by 312 data segments, a second data field sync segment, another 312 data segments.

Except for the binary data segment and data field syncs, all other transmitted symbols are multi-level. For the terrestrial broadcast mode, 8-level symbols (3 bits/symbol) are transmitted while for the high data rate cable mode, 16-level symbols (4 bits/symbol) are used. These are called 8-VSB and 16-VSB, respectively.

The multi-level symbols combined with the data segment and data field syncs are used to suppress-carrier modulate a single carrier. However, before transmission, the lower sideband is removed. The resulting spectrum is flat, except for the band edges where a root-raised cosine response results in 620 KHz transition regions. The VSB transmission spectrum relative to an NTSC spectrum is shown in Fig 2. The cumulative distribution function (CDF) of the peak-to-average power ratio of a typical transmitted signal is plotted in Fig 3, with 99.9% of the power envelope within 6.3 dB of the average power.

At the suppressed carrier frequency of 310 KHz from the lower band edge, a small pilot is also added to the signal. The pilot is used in the VSB receiver to achieve carrier lock. Pilot power adds 0.3 dB to the total signal power but helps reduce implementation loss by more than that. With the aid of the pilot, the VSB transmission system achieves virtual theoretical performance (i.e. no implementation loss). The pilot is positioned in the vestigial sideband region of a cochannel NTSC signal and does not contribute to cochannel interference into NTSC.

The terrestrial VSB system was designed with robustness in mind. Forward error correction in the form of R-S and trellis coding, along with 1/3 data field interleaving, provide a rugged system that can endure both white noise and interference environments. The terrestrial VSB system can operate in signal-to-noise (S/N) environments of 14.9 dB ($3x10^{-6}$), as illustrated in the 8-VSB, 4-state segment error probability curve in Fig 4.

6.2.2 Transmitter Broadcast Mode

A functional block diagram of the VSB terrestrial broadcast transmitter is shown in Fig 5. Descriptions of each block follow.

6.2.2.1 Data Randomizer

A data randomizer is used on all input data to randomize the data payload (NOT including syncs, or R-S parity bytes). This ensures random data is transmitted even when constant data is applied to the system, as might happen when a data input is disconnected. Random data is important for optimum reception of the transmitter data signal.

The data randomizer exclusive OR's all the incoming data bytes with a 16-bit maximum length pseudo random sequence (PRS) locked to the data frame. The PRS is generated in a 16-bit shift register that has 9 feedback taps. Eight of the shift register outputs are selected as the fixed randomizing byte, where each bit from this byte is used to individually exclusive OR the corresponding input data bit.

6.2.2.2 Reed-Solomon Encoder

The Reed-Solomon (R-S) error protection code, known for its burst noise correction capability, is a very efficient block code for its parity overhead. Since it works with bytes (8-bits), any number of bit errors can occur within a byte and not affect its error correction ability. Therefore, there is no direct proportional relationship between byte error rate and bit error rate. Only at very low error rates do the two converge, but the bit error rate is never smaller than the byte error rate for the same S/N ratio. The measure of performance for block codes is its packet error rate versus signal-to-noise ratio or signal-to-interference ratio. Since ATV data is to be transmitted in MPEG transport packets, the most common method of describing data transmission performance in noise and interference conditions is probability of data packet error.

The R-S code used in the VSB transmission system is a t=10 (208,188) code. The data block size is 188 bytes, with 20 R-S parity bytes added for error correction. A total block size of 208 bytes is transmitted per data segment. A t=10 R-S code, with 20 parity bytes, can correct up to 10 byte errors per block.

6.2.2.3 Interleaver

Although the R-S code is particularly powerful in protecting against burst errors, the data is interleaved for further protection. Burst errors result from both impairments and trellis decoding mistakes where a single error can propagate through the trellis decoder, multiply, and become a burst error. Long trellis decoders (large K) cause correspondingly longer burst errors to be created.

The goal of the interleaver is to spread the data bytes from the same R-S block over time so that a long burst of noise or interference is necessary to encompass more than 10 data bytes and overrun the R-S error protection. Since the R-S code can correct up to 10 byte errors per block regardless of how many individual bit errors have occurred within each byte, it is advantageous to group consecutive symbols from the trellis decoder into bytes to get the best burst error performance.

The interleaver employed in the VSB transmission system is an 87 data segment (intersegment) diagonal byte interleaver. Interleaving is provided to a depth of about 1/3 of a data field (5.5 msec deep). Intrasegment interleaving is also performed for the benefit of the trellis coding process, to be described later.

6.2.2.4 Trellis Coded Modulation

The terrestrial broadcast VSB transmission mode employs a 2/3 rate (R=2/3) trellis code (with one unencoded bit). That is, one input bit is encoded into two output bits using a 1/2 rate convolution code while the other input bit is left unencoded. The signaling waveform used with the trellis code is an 8-level (3 bit) 1-dimensional constellation. The transmitted signal is commonly referred to as 8-VSB. To minimize the complexity of interleaving and trellis decoder hardware, as well as to minimize error propagation, a relatively simple (short) 4-state trellis encoder is used. Long trellis codes, which cause longer burst errors and require more interleaving, were not found to be beneficial.

Although trellis codes produce improvements in signal-to-noise ratio (S/N) threshold against white noise, they do not perform well for impulse or burst noise. Besides electromechanical sources of burst noise, burst noise is also caused by NTSC cochannel interference and phase noise which can cause data-dependent cross talk. To further reduce the effects of burst errors, and to simplify the trellis decoder when the NTSC rejection filter is used, trellis code interleaving is used. This uses twelve identical trellis encoders operating on interleaved data symbols. The code interleaving is accomplished by encoding symbols {0,12,24,36...} as one group, symbols {1,13,25,37} as a second group, symbols {2,14,26,38,...} as a third group, and so on for a total of 12 groups. The 12 groups are required since the creation of spectral nulls in the receiver requires a comb filter with delays of 12 symbols. This will be discussed further later.

The theoretical performance of the concatenated trellis and R-S code used in the terrestrial broadcast VSB system was shown earlier in Fig 4. A theoretical performance curve for the 4-VSB system as delivered for the first round of ATV testing is also shown for reference. The combination of short trellis codes and code interleaving is a superior error correction technique for use in channels with white gaussian noise and Rayleigh distribution fading.

6.2.2.5 Synchronization Information

Synchronization information is added to the digital data signal in order to facilitate packet and symbol clock acquisition and phase- lock during extreme noise and interference conditions. The encoded trellis data is passed through a multiplexer that inserts the various synchronization signals (data segment sync & data frame sync).

A two level (binary) 4-symbol Data Segment Sync is inserted into the 8-level digital data stream at the beginning of each data segment. The data segment sync embedded in random data is illustrated in Fig 6. A complete segment consists of 836 symbols: 4 sync symbols, and 832 data plus parity symbols. The segment sync is binary (2-level) in order to make packet and clock recovery rugged. The levels selected (±5) for segment sync insure robustness, but do not cause undue interference into cochannel NTSC signals. They also have an average value of zero so that when the DC pilot is added, they will not alter the desired value of pilot. The same sync pattern occurs regularly at 77.695 msec intervals, and are the only signals repeating at this rate. The periodic recurrence of these four symbols makes possible their reliable detection at the receiver under severe noise and/or interference. Unlike the data, the four data segment sync symbols are not Reed-Solomon or trellis encoded, nor are they interleaved.

The data is not only divided into data segments, but also into data fields. Each data field starts with one complete data segment of Data Field Sync, as shown in Fig 7. The Data Field Sync signal is made up of binary (two-level) pseudo-random sequences of 511 and 63 symbols. The 12 symbols at the end of the reference signal are repeated from the previous active data segment to aid the trellis decoder in the receiver during cochannel interference when the NTSC rejection filter (12 symbol subtractive comb) is in the data path. Binary levels provide reliable detection for the Data Field sync under extremely adverse channel conditions.

The three concatenated 63 bit PRBS sequences in the field sync alternate polarity from one field to the next, providing field identification information to the receiver. Therefore, there are two data field syncs that make up one data frame (48.6 msec). The average value of the field sync symbols is close to zero so that when the DC pilot is added prior to transmission, they will not alter the desired value of pilot.

The Data Field Sync serves five purposes. First, it provides a means to determine the beginning of each data field. Second, it is also used by the equalizer in the ATV receiver as a training reference signal to remove intersymbol and other interferences. Third, it allows the receiver to determine whether the interference rejection filter (12 symbol subtractive comb) should be used. Fourth, it is used for system diagnostic measurements, such as signal-to-noise and channel response. Fifth, the phase tracker in the receiver uses the field sync to reset its circuitry and determine its loop parameters. Just like the data segment sync, the Data Field Sync is not Reed-Solomon or trellis encoded, nor is it interleaved.

6.2.2.6 Pilot Insertion

A rugged system must be able to acquire a signal and maintain lock in the presence of very heavy noise and interference. A small pilot added to the suppressed carrier RF data signal is adequate to allow robust carrier recovery in the receiver during these extreme conditions.

A small (digital) DC level (1.25) is added to every symbol (data and syncs) of the digital baseband data plus sync signal (±1, ±3, ±5, ±7). This has the effect of adding a small in-phase pilot to the data signal. Digital addition of the pilot (at baseband) provides a highly stable and accurate pilot. The frequency of the pilot will be the same as the suppressed carrier frequency. Since the data is essentially guaranteed to be random over long intervals, all data states are equally probable. Using the above eight values for the random data symbols, the total average data power of the 8-level data signal is 21. After adding the pilot, the total average signal power is 22.56. Thus, the pilot represents an increased transmitted power of 0.3 dB. In the interference-limited environment, the pilot is not a contributing factor. Also, the pilot power has no significant effect on the ATV transmitter hardware (e.g. power dissipation or peak-to-average power ratio).

6.2.2.7 Pre-Equalizer Filter

A pre-equalizer filter is available for use in over-the-air broadcasts, where the high power transmitter may have significant in-band ripple or roll off at band edges. This linear distortion can be detected by an equalizer in a reference demodulator ("ideal" receiver) located at the transmitter site that is receiving a small portion of the antenna signal feed (from a "sample"). The reference demodulator equalizer tap weights can be transferred into the transmitter pre-equalizer for pre-correction of transmitter linear distortion.

The pre-equalizer is an 80 tap, feedforward transversal filter. It operates on the I channel data signal (there is no Q channel data in the transmitter), and shapes the frequency spectrum of the IF signal so that it is a flat spectrum at the output of the high power transmitter that feeds the antenna for transmission.

6.2.2.8 Vsb Modulator

The VSB modulator receives the 10.76 Msymbols/sec, 8-level trellis composite data signal (pilot plus syncs added). For minimal intersymbol interference, the data signal must be properly shaped (filtered) before it is transmitted over the 6 MHz channel. A linear phase raised-cosine Nyquist filter is employed in the concatenated transmitter and receiver, as shown in Fig 8. The system filter response is essentially flat across the entire band, except for the transition regions at each end of the band. Due to the vestigial-sideband nature of the transmitted signal, the same skirt selectivity on both sides is not required, although it is so implemented. The optimum system arrangement is to divide the roll off equally between the transmitter and receiver filters. Therefore, root-raised cosine filters are used. Fig 2 illustrates the ATV transmitter spectrum in comparison with a typical NTSC spectrum.

The transmitter VSB filtering is implemented by complex-filtering the baseband data signal, creating precision-filtered and stable in- phase and quadrature-phase modulation signals. This filtering process provides the root-raised cosine Nyquist filtering as well as the sin x/x compensation for the D/A converters. The orthogonal baseband signals are converted to analog form (D/A converters) and then modulated on quadrature IF carriers to create the vestigial sideband IF signal by sideband cancellation (phasing method). The nominal frequency of the IF carrier (and small in-phase pilot) is 46.69 MHz, which is equal to the IF center frequency (44.000 MHz) plus the data rate divided by 4 (10.762 MHz/4=2.6905 MHz). Additional adjacent channel suppression (beyond that achieved by sideband cancellation) is performed by a linear phase, flat amplitude response SAW filter. Adjacent channel energy spillage at the IF output is at least 57 dB down from the desired ATV signal power.

6.2.2.9 Upconverter And Rf Carrier Frequency Offsets

Modern NTSC TV transmitters use a two-step modulation process. The first step usually is modulation of the data onto an IF carrier, which is the same frequency for all channels, followed by translation to the desired RF channel. The VSB transmitter applies this same two-step modulation process. The RF upconverter translates the filtered flat IF data signal spectrum to the desired RF channel. For the same coverage as an NTSC transmitter, the average power of the ATV signal will be at least 12 dB less than the NTSC peak sync power.

The frequency of the RF upconverter oscillator in ATV terrestrial broadcasts will typically be the same as that used for NTSC (except for NTSC offsets). However, in extreme cochannel situations, the ATV system is designed to take advantage of precise RF carrier frequency offsets with respect to the NTSC cochannel carrier. Since the VSB data signal sends repetitive synchronizing information (segment syncs), precise offset causes NTSC cochannel carrier interference into the VSB receiver to phase alternate from sync to sync. The VSB receiver circuits average successive syncs to cancel the interference and make data segment sync detection more reliable.

For ATV cochannel interference into NTSC, the interference is noise-like and does not change with precise offset. Even the ATV pilot interference into NTSC does not benefit from precise frequency offset because it is so small (11.3 dB below the data power) and falls far down the Nyquist slope (20 dB or more) of NTSC receivers.

Although it might be postulated that an ATV transmitter can be located so as to experience equal interference from two worst-case cochannel NTSC stations (e.g. three-way triangle), such a situation is so unlikely that the ATV signal is assumed to have only one dominant NTSC cochannel. The ATV cochannel pilot should be offset in the RF upconverter from the dominant NTSC picture carrier by an odd multiple half the data segment rate. A $56.9*f_H$ = 895.28 KHz offset is proposed between the pilot frequency and that of the cochannel NTSC picture carrier. In order to achieve this requirement, a 45.8 KHz spectrum shift of the VSB signal into the upper adjacent channel is required. This is illustrated (Fig 15) and further described later in the section on the receiver NTSC rejection filter.

For ATV to ATV cochannel interference, precise carrier offset prevents possible misconvergence of the adaptive equalizer. If perchance the two ATV Data Field Sync signals should fall within the same data segment time, the adaptive equalizer could misinterpret the interference as a ghost. To prevent this, a carrier offset of $f_{seg}/2$ = 6.437 KHz is proposed for close ATV-to-ATV cochannel situations. This causes the interference to have no effect in the adaptive equalizer.

6.2.3 Receiver Broadcast Mode

Fig 9 shows the receiver block diagram of the VSB terrestrial broadcast transmission system. Descriptions of each block follow.

6.2.3.1 Tuner

The tuner, illustrated in Fig 10, receives the 6 MHz ATV signal from the antenna. It is a high-side injection double-conversion type with a first IF frequency of 920 MHz. This puts the image frequencies above 1 GHz, making them easy to reject by a fixed front end filter. This selection of first IF frequency is high enough so that the input bandpass filter selectivity prevents the local oscillator (978-1723 MHz) from leaking out the tuner front end and interfering with other UHF channels, yet it is low enough for second harmonics of UHF channels (470-806 MHz) to fall above the first IF bandpass. Harmonics of cable channels could possibly occur in the first IF passband but are not a real problem because of the relatively flat spectrum (within 10 dB) and small signal levels (-28 dBm or less) used in cable systems.

The tuner input has a bandpass filter that limits the frequency range to 50-810 MHz, rejecting all other non-television signals that may fall within the tuners image frequency range (beyond 920 MHz). In addition, a broadband tracking filter rejects other television signals, especially those much larger in signal power than the desired ATV signal power. This tracking filter is not narrow, nor is it critically tuned, as is the case of present day NTSC tuners that must reject image signals only 90 MHz away from the desired channel. Minimal channel tilt, if any, exists due to this tracking filter.

A 10 dB gain, wideband RF amplifier increases the signal level into the first mixer, and is the dominant determining factor of receiver noise figure (7-9 dB over entire VHF, UHF, and cable bands). The first mixer, a highly linear double-balanced design to minimize even harmonic generation, is driven by a synthesized low phase noise local oscillator (LO) above the first IF frequency (high-side injection). Both the channel tuning (first LO) and broadband tracking filtering (input bandpass filter) are controlled by microprocessor. The tuner is capable of tuning the entire VHF and UHF broadcast bands as well as all standard, IRC, and HRC cable bands.

The mixer is followed by an LC filter in tandem with a narrow 920 MHz bandpass ceramic resonator filter. The LC filter provides selectivity against the harmonic and subharmonic spurious responses of the ceramic resonators. The 920 MHz ceramic resonator bandpass filter has a -1 dB bandwidth of about 6 MHz. A 920 MHz IF amplifier is placed between the two filters. Delayed

AGC of the first IF signal is applied immediately following the first LC filter. The 30 dB range AGC circuit protects the remaining active stages from large signal overload.

The second mixer is driven by the second LO, which is an 876 MHz voltage-controlled SAW oscillator. It is controlled by the frequency and phase-locked loop (FPLL) synchronous detector. The second mixer, whose output is the desired 44 MHz second IF frequency, drives a constant gain 44 MHz amplifier. The output of the tuner feeds the IF SAW filter and synchronous detection circuitry.

The tuner is made out of standard consumer electronic components, and is housed in a stamped metal enclosure.

6.2.3.2 Channel Filtering And Vsb Carrier Recovery

Carrier recovery is performed on the small pilot carrier by an FPLL circuit, illustrated in Fig 11. The first LO is synthesized by a PLL and controlled by a microprocessor. The third LO is a fixed reference oscillator. Any frequency drift or deviation from nominal has to be compensated in the second LO. Control for the second LO comes from the FPLL synchronous detector, which integrally contains both a frequency loop and a phase-locked loop in one circuit. The frequency loop provides a wide frequency pull-in range of ±100 KHz while the phase-locked loop has a narrow bandwidth (less than 2 KHz).

During frequency acquisition, the frequency loop uses both the in-phase (I) and quadrature-phase (Q) pilot signals. All other data processing circuits in the receiver use only the I-channel signal. Prior to phase-lock, as is the condition after a channel change, the automatic frequency control (AFC) lowpass filter acts on the beat signal created by the frequency difference between the VCO and the incoming pilot. The high frequency data (as well as noise and interference) is mostly rejected by the AFC filter, leaving only the pilot beat frequency. After limiting this pilot beat signal to a constant amplitude (±1) square wave, and using it to multiply the quadrature signal, a traditional bipolar S-curve AFC characteristic is obtained. The polarity of the curve depends upon whether the VCO frequency is above or below the incoming IF signal. Filtered and integrated by the automatic phase control (APC) lowpass filter, this DC signal adjusts the tuner's second LO to reduce the frequency difference.

When the frequency difference comes close to zero, the APC loop takes over and phase-locks the incoming IF signal to the third LO. This is a normal phase-locked loop circuit, with the exception that it is bi-phase stable. However, the correct phase-lock polarity is determined by forcing the polarity of the pilot to be equal to the known transmitted positive polarity. Once locked, the detected pilot signal is constant, the limiter output feeding the third multiplier is at a constant +1, and only the phase-locked loop is active (frequency loop automatically disabled). The APC lowpass filter is wide enough to reliably allow ±100 KHz frequency pull-in, yet narrow enough to consistently reject all strong white noise (including data) and NTSC cochannel interference signals. The PLL has a bandwidth that is narrow enough to reject most of the AM and PM generated by the data, yet is wide enough to track out any phase noise on the signal (and, hence, on the pilot) out to about 2 KHz. Tracking out low frequency phase noise (as well as low frequency FM components) allows the phase tracking loop, discussed later, to be more effective.

The prototype receiver can acquire a signal and maintain lock at a signal-to-noise of 0 dB or less, and in the presence of heavy interference.

6.2.3.3 Segment Sync And Symbol Clock Recovery

The repetitive data segment syncs are detected from among the synchronously detected random data by a narrow bandwidth filter. From the data segment syncs, a properly phased 10.76 MHz symbol clock is created along with a coherent AGC control signal. A block diagram of this circuit is shown in Fig 12.

The 10.76 Msymbols/sec I-channel composite baseband data signal (syncs and data) from the synchronous detector is converted by an A/D converter for digital processing. Traditional analog data eyes can be viewed after synchronous detection. However, after conversion to a digital signal, the data eyes cannot be seen due to the sampling process. A PLL is used to derive a clean 10.76 MHz symbol clock for the receiver.

With the PLL free-running, the data segment sync detector containing a 4-symbol sync correlator looks for the two level syncs occurring at the specified repetition rate. The repetitive segment sync is detected while the random data is not, enabling the PLL to lock on the sampled sync from the A/D converter, and achieve data symbol clock synchronization. Upon reaching a predefined level of confidence (using a confidence counter) that the segment sync has been found, subsequent receiver loops are enabled.

Data segment sync detection and clock recovery both work reliably at signal-to-noise ratios of 0 dB or less, and in the presence of heavy interference.

6.2.3.4 Non-Coherent And Coherent Agc

Prior to carrier and clock synchronization, non-coherent automatic gain control (AGC) is performed whenever any signal (locked or unlocked signal, or noise/interference) overruns the A/D converter. The IF and RF gains are reduced accordingly, with the appropriate AGC "delay" applied.

When data segment syncs are detected, coherent AGC occurs using the measured segment sync amplitudes. The amplitude of the bipolar syncs, relative to the discrete levels of the random data, is determined in the transmitter. Once the syncs are detected in the receiver, they are compared to a reference value, with the difference (error) integrated. The integrator output then controls the IF and "delayed" RF gains, forcing them to whatever values provide the correct sync amplitudes.

6.2.3.5 Data Field Synchronization

Data Field Sync detection, shown in Fig 14, is achieved by comparing each received data segment from the A/D converter (after interference rejection filtering to minimize cochannel interference), to ideal field #1 and field #2 reference signals in the receiver. Oversampling of the field sync is NOT necessary since a precision data segment and symbol clock has already been reliably created by the clock recovery circuit. Therefore, the field sync recovery circuit knows exactly where a valid field sync correlation should occur within each data segment, and only needs to perform a symbol by symbol difference. Upon reaching a predetermined level of confidence (using a confidence counter) that field syncs have been detected on given data segments, the Data Field Sync signal becomes available for use by subsequent circuits. The polarity of the three alternating 63 bit pseudo random (PN) sequences determine whether field 1 or field 2 is detected.

This procedure makes field sync detection robust, even in heavy noise, interference, or ghost conditions. Field sync recovery can reliably occur at signal-to-noise ratios of 0 dB or less, and in the presence of heavy interference.

6.2.3.6 Interference Rejection Filter

The interference rejection properties of the VSB transmission system are based on the frequency location of the principal components of the NTSC cochannel interfering signal within the 6 MHz TV channel and the periodic nulls of a VSB receiver baseband comb filter.

Fig 15a shows the location and approximate magnitude of the three principal NTSC components: (1) the visual carrier (V) located 1.25 MHz from the lower band edge, (2) the chrominance subcarrier (C) located 3.58 MHz higher than the visual carrier frequency, and (3) the aural carrier (A) located 4.5 MHz higher than the visual carrier frequency.

The NTSC interference rejection filter (comb) is a one tap linear feed-forward filter, as shown in Fig 16. Fig 15b shows the frequency response of the comb filter, which provides periodic spectral nulls spaced 57* f_h (10.762 MHz/12, or 896.85 KHz) apart. There are 7 nulls within the 6 MHz channel. The NTSC visual carrier frequency falls close to the second null from the lower band edge. The 6th null from the lower band edge is correctly placed for the NTSC chrominance subcarrier, and the 7th null from the lower band edge is near the NTSC aural carrier.

Comparing Fig 15a and Fig 15b shows that the visual carrier falls 2.1 kHz below the second comb filter null, the chroma subcarrier falls exactly at the 6th null, and the aural carrier falls into the 7th null. (Note, the aural carrier is at least 7 dB below its visual carrier).

The comb filter, while providing rejection of steady-state signals located at the null frequencies, has a finite response time of 12 symbols (1.115 msecs). Thus, if the NTSC interfering signal has a sudden step in carrier level (low to high or high to low), one cycle of the beat frequency (offset) between the ATV and NTSC carrier frequencies will pass through the comb filter at an amplitude proportional to the NTSC step size as instantaneous interference. Examples of such steps of NTSC carrier are: leading and trailing edge of sync (40 IRE units). If the desired to undesired (D/U) signal power ratio is large enough, data slicing errors will occur. However, interleaving will spread the interference and will make it easier for the Reed-Solomon code to correct them (R-S can correct up to 10 byte errors/segment).

Although the comb filter reduces the NTSC interference, the data is also modified. The 7 data eyes (8 levels) are converted to 14 data eyes (15 levels). Such a process is called partial response. The modified data signal can be properly decoded by the trellis decoder, and will be described in later sections. Note, because of time sampling, only the maximum data eye value is seen after A/D conversion.

The detail at the band edges for the overall channel is shown in Fig 15c and Fig 15d. Fig 15d shows that the frequency relationship (56«INVALID_FIELD: Object»*f_H between NTSC visual carrier and ATV carrier) requires a shift in the ATV spectrum with respect to the nominal channel. The shift equals +45.8 KHz, or about +0.76%. This is slightly higher than currently applied channel offsets and reaches into the upper adjacent channel at a level of about -40 dB. If that is another ATV channel, the nominal situation is restored. If it is an NTSC channel, the shift is below the (RF equivalent of the) Nyquist slope of an NTSC receiver where there is high attenuation, and it is slightly above its customary lower adjacent channel sound trap. No adverse effects of the shift have been found nor are they foreseen.

NTSC interference can be detected by the circuit shown in Fig 16, where the signal-to-interference plus noise ratio of the binary Data Field Sync is measured at the input and output of the comb filter, and compared to each other. This is accomplished by creating an error signal at each rejection filter port by comparing the received signal with a stored reference of the field sync. The errors are squared and integrated. After a predetermined level of confidence is achieved, the path with the largest signal-to-noise ratio (lowest interference energy) is switched in and out of the system automatically.

There is a reason to not leave the rejection comb filter switched in all the time. The comb filter, while providing needed cochannel interference benefits, degrades white noise performance by 3 dB. This is due to the fact that the filter output is the subtraction of two full gain paths, and since white noise is uncorrelated from symbol to symbol, the noise power doubles. If little or no NTSC interference is present, the comb filter is automatically switched out of the data path.

6.2.3.7 Channel Equalizer

The equalizer/ghost canceler is used to compensate for linear channel distortions, such as tilt and ghosts. The distortions can come from the transmission channel or from imperfect components within the receiver. The equalizer delivered for testing uses a Least-Mean-Square (LMS) algorithm and adapts on the transmitted binary Data Field Sync as well as on the random data. A block diagram of the equalizer used for the G-A "bakeoff" testing is shown in Fig 17.

The advantage of adapting on a known training signal embedded within the random data signal is the guarantee of tap convergence, especially in extreme noise, interference, and ghost conditions. All the various loops within the receiver (carrier, clock, AGC, segment sync, frame sync) are independent of subsequent loops, including the equalizer. An optimum method is to use a two step process: equalization first on a binary training sequence to open the data eyes, and then equalization on the random data for high speed tracking of moving ghosts (e.g. airplane flutter).

The equalizer filter consists of two parts, a 78-tap feedforward transversal filter followed by a 177-tap decision-feedback section. The long decision-feedback provides more accurate ghost canceling due to reduced noise enhancement than that found in long feedforward equalizers. The equalizer operates at the 10.762 MHz symbol rate (T-sampled equalizer). Symbol rate sampling is made possible because the symbol clock has been accurately phased by the clock recovery circuitry BEFORE the equalizer. To aid convergence, the equalizer also includes an extra adder which is used to add or subtract a DC value to compensate for DC errors which can be caused by circuit offsets, nonlinearities, or shifts in the pilot due to ghosts.

For the final G-A system, the G-A intends to augment the equalizer with a blind equalization algorithm to improve airplane flutter performance.

6.2.3.8 Phase Tracking Loop

The phase tracking loop is an additional decision feedback loop which further tracks out phase noise which has not been removed by the IF PLL operating on the pilot. Thus, phase noise is tracked out by not just one loop, but two concatenated loops. Because the system is already frequency-locked to the pilot by the IF PLL (independent of the data), the phase tracking loop bandwidth is maximized for phase tracking by using a first order loop. Higher order loops, which are needed for frequency tracking, do not perform phase tracking as well as first order loops. Therefore, they are not used in the VSB system.

A block diagram of the phase tracking loop is shown in Fig 18. The output of the real equalizer operating on the I signal is first gain controlled by a multiplier and then fed into a filter which recreates an approximation of the Q signal. This is possible because of the VSB transmission method, where the I and Q components are related by a filter function which is almost a Hilbert transform. The complexity of this filter is minor since it is a finite impulse response (FIR) filter with fixed anti-symmetric coefficients and with every other coefficient equal to zero. In addition, many filter coefficients are related by powers of two, thus simplifying the hardware design.

These I and Q signals are then fed into a de-rotator (complex multiplier), which is used to remove the phase noise. The amount of de-rotation is controlled by decision feedback of the data taken from the output of the de-rotator. Since the phase tracker is operating on the 10.76 Msymbol/sec data, the bandwidth of the phase tracking loop is fairly large, approximately 60 KHz. The gain multiplier is also controlled with decision feedback.

6.2.3.9 Trellis Decoder

To help protect the trellis decoder against short burst interference, such as impulse noise or NTSC cochannel interference, 12 symbol code interleaving is employed in the transmitter. As shown in Fig 19, the receiver uses 12 trellis decoders in parallel, where each trellis decoder sees every 12th symbol. This code interleaving has all the same burst noise benefits of a 12 symbol interleaver, but also minimizes the resulting code expansion (and hardware) when the NTSC rejection comb filter is active.

The trellis decoder performs the task of slicing and convolutional decoding. It has two modes; one when the NTSC rejection filter is used to minimize NTSC cochannel, and the other when it is not used. This is illustrated in Fig 20. The insertion of the NTSC rejection filter is determined automatically (before the equalizer), with this information passed to the trellis decoder. When there is little or no NTSC cochannel interference, the NTSC rejection filter is not used, and an optimal trellis decoder is used to decode the 4-state trellis-encoded data.

In the presence of significant NTSC cochannel interference, when the NTSC rejection filter (12 symbol, feedforward subtractive comb) is employed, a trellis decoder optimized for this partial response channel is used. This optimal code requires 8 states. This is necessary since the NTSC rejection filter, which has memory, represents another state machine seen at the input of the trellis decoder. In order to minimize the expansion of trellis states, two measures are taken: (1) special design of the trellis code, and (2) twelve-to-one interleaving of the trellis encoding. The interleaving, which corresponds exactly to the 12 symbol delay in the NTSC rejection filter, makes it so that each trellis decoder only sees a one-symbol delay NTSC rejection filter. By minimizing the delay stages seen by each trellis decoder, the expansion of states is also minimized. Only a 3.5 dB penalty in white noise performance is paid as the price for having good NTSC cochannel performance. The additional 0.5 dB beyond the 3 dB comb filter noise threshold degradation is due to the 12 symbol differential coding.

It should be noted that after the ATV transition period and NTSC is no longer being transmitted, the NTSC rejection filter and the 8-state trellis decoder can be eliminated from receivers.

Trellis decoder complexity is exponentially proportional to the number of states. The 4-state and 8-state decoders used in the VSB terrestrial broadcast mode are very simple and cost effective, yet result in about 1.50 dB improvement in signal-to-noise ratio over the 4-VSB system. Increasingly complex trellis codes, which provide small amounts of theoretical improvement in coding gain, will require increased interleaving complexity, and are also more susceptible to burst errors and implementation loss.

6.2.3.10 Data De-Interleaver

The de-interleaver performs the exact inverse function of the transmitter interleaver. Its 1/3 data field depth, and intersegment "dispersion" properties allow noise bursts lasting about 170 msecs to be handled. Even strong NTSC cochannel signals passing through the NTSC rejection filter and creating short bursts due to NTSC vertical edges, are reliably handled due to the interleaving and R-S coding process. The intrasegment interleaving is also performed, reassembling the trellis data from the 12 trellis decoders into their original form.

6.2.3.11 Reed-Solomon Decoder

The trellis-decoded byte data is sent to the (208,188) t=10 R-S decoder, where it uses the 20 parity bytes to perform the byte-error correction on a segment-by-segment basis. Up to 10-byte errors/data segment are corrected by the R-S decoder. Any burst errors created by impulse noise, NTSC cochannel interference, or trellis-decoding errors, are greatly reduced by the combination of the interleaving and R-S error correction.

6.2.3.12 Data De-Randomizer

The de-randomizer accepts the error-corrected data bytes from the R-S decoder, and applies the same PRS randomizing code to the data. The PRS code is generated identically as in the transmitter, using the same PRS generator feedback and output taps. Since the PRS is locked to the reliably recovered Data Field Sync (and not some codeword embedded within the potentially noisy data), it is exactly synchronized with the data, and performs reliably.

6.2.3.13 Receiver Loop Acquisition Sequencing

The receiver incorporates a "universal reset" which initiates a number of "confidence counters" and "confidence flags" involved in the lock-up process. A universal reset occurs, for example, when tuning to another station or turning on the receiver.

The various loops within the VSB receiver acquire and lock-up sequentially, with "earlier" loops being independent from "later" loops. The order of loop acquisition is as follows:

* Tuner 1st LO synthesizer acquisition
* Non-coherent AGC reduces unlocked signal to within A/D range
* Carrier acquisition (FPLL)
* Data segment sync and clock acquisition
* Coherent AGC of signal (IF and RF gains properly set)
* Data field sync acquisition
* NTSC rejection filter insertion decision made
* Equalizer completes tap adjustment algorithm
* Trellis and R-S data decoding begin